arXiv:0712.0483v1 [quant-ph] 4 Dec 2007

Interacting electrons, Density Functional Theory, and Quantum Merlin Arthur

Norbert Schuch[1] and Frank Verstraete[2]

[1] Max-Planck-Institut für Quantenoptik,
Hans-Kopfermann-Str. 1, D-85748 Garching, Germany.
email: `norbert.schuch@mpq.mpg.de`

[2] Fakultät für Physik, Universität Wien,
Boltzmanngasse 5, A-1090 Wien, Austria.
email: `frank.verstraete@univie.ac.at`

Abstract

During the last decades, the formalism of computational complexity has played a central role in advancing the understanding of the physics of spin glasses and in explaining the difficulty of simulating problems in statistical physics. The advent of quantum complexity theory has put us into the unique position to foster a similar understanding for systems of interacting quantum particles. The central problem in quantum mechanics, and more specifically in quantum chemistry and solid state physics, is to solve the Schrödinger equation for a system of electrons interacting via the Coulomb potential in an external field. We show that this problem is QMA (Quantum Merlin Arthur) hard by reducing it to a 2D lattice of qubits with a fixed Heisenberg interaction and local magnetic fields, which in turn can be reduced to a known QMA-complete problem. On our way, we find that the Hubbard model in the presence of local magnetic fields, one of the most intensively studied lattice models for fermions, is QMA-complete.

This result is highly relevant within the context of Density Functional Theory (DFT), the most widely used and successful method for simulating systems of interacting electrons. In a series of papers that seem to be the most cited ones ever in physics, Kohn and collaborators proved the existence of a problem-independent black box universal functional [5] whose usage would considerably simplify the simulation of interacting electrons. The central problem in DFT has been to find better and better approximations to this universal functional. We show that the existence of good approximations would place the problem of interacting electrons in NP, implying NP=QMA. This provides strong evidence against the existence of an efficiently computable universal functional and poses bounds on its approximability. We finally discuss another widely used simulation method for interacting electrons, the Hartree-Fock ansatz, and show that finding the optimal energy there is NP-complete.

1 Motivation and problem

In this paper, we consider the problem of finding the ground state of a system of interacting electrons moving in an external potential from the point of view of computational complexity. This problem is encountered virtually everywhere in quantum chemistry as well as in condensed matter physics. For instance, the spatial configuration of a molecule is the one for which the energy of the interacting electrons moving in the nuclear potential, together with the electrostatic energy of the nuclei, gets minimal. Similarly, a rich variety of phenomena in solid state physics, in particular conductance and magnetic phenomena, can be understood by considering electrons moving in the periodic lattice potential, including such exciting phenomena as high-temperature superconductivity and the fractional quantum Hall effect.

Let us first define the problem under consideration. We are given a system of N electrons, which is described by a complex-valued $\mathcal{L}_2$ function $\Psi(x_1,\dots,x_N)$, where each $x_i=(r_i,s_i)$ encodes the spatial coordinate $r_i\in\mathbb{R}^3$ and the spin $s_i\in\{0,1\}$, and where the fermionic nature of the electrons requires Ψ to be antisymmetric,

$$\Psi(\dots,x_i,\dots,x_j,\dots)=-\Psi(\dots,x_j,\dots,x_i,\dots)$$

for any pair of indices i and j. On the space of completely antisymmetric states, we want to find the ground state energy – the lowest eigenvalue – of the Hamiltonian

$$H=\underbrace{-\tfrac{1}{2}\sum_i\Delta_i}_{=:T}+\underbrace{\sum_{i<j}\frac{1}{|r_i-r_j|}}_{=:I}+\sum_i V(x_i)\ , \tag{1}$$

where the external potential V contains an electrostatic part $\phi(r_i)$ and a magnetic part $\vec{B}(r_i)$, such that $V(r_i,s_i)=\phi(r_i)-\vec{B}(r_i)\cdot\vec{\sigma}_i$ with the Pauli matrices $\vec{\sigma}_i=(\sigma_i^x,\sigma_i^y,\sigma_i^z)$ acting on the spin degree of freedom, and where the Laplace operator Δ_i is with respect to the i'th spatial coordinate.

The computational task we consider is finding the ground state energy of this system, up to a polynomial precision in energy. More precisely, we will focus on the corresponding decision problem INTERACTING ELECTRONS, where we have to decide whether the ground state energy is below some value a or above b, where $b-a>1/\mathrm{poly}(N)$ is the desired precision.

The natural complexity class for ground state problems is QMA, Quantum Merlin Arthur [10]. A language L is in QMA if there exists a uniform family of polynomial-time quantum circuits V_x such that for every $x\in L$, there is a polynomial-size quantum state $|\psi\rangle$ (the *proof*) such that V_x will accept $|\psi\rangle$ with probability at least $2/3$, while for $x\notin L$, V_x will accept any proof with probability at most $1/3$. The canonical complete problem for the class QMA is LOCAL HAMILTONIAN [10], where a local Hamiltonian on N qubits is given and the problem is to decide whether its ground state energy is below a or above b for $b-a>1/\mathrm{poly}(N)$. Here, by local we mean interactions involving at most a constant number of qubits each. Subsequently, there have been found more restricted ground state problems which are still QMA-complete, most notably for qubits on a 2D square lattice with a particular class of nearest neighbor couplings [13] and a 1D chain of 12-level systems with nearest neighbor couplings [1].

In this article, we determine the computational complexity of INTERACTING ELECTRONS, which we show to be QMA-hard. Therefore, we consider the Hubbard model which describes electrons hopping on a lattice and is well-known to be a limiting case of the general problem of interaction electrons, Eq. (1); this reduces the original problem, defined on a continuous system, to a lattice problem for fermions. In a second step, we reduce the fermionic system to a system of distingushable spins, i.e. qubits, by showing that there is a domain in which the Hubbard model mimics the behaviour of qubits coupled via a fixed-strength Heisenberg interaction and subject to local magnetic fields. We then show how this Hamiltonian can be reduced to a model which is known to be QMA-complete, namely a two-dimensional square lattice of qubits coupled via so-called Pauli interactions [13].

Having shown QMA-completeness of INTERACTING ELECTRONS, we turn our attention to Density Functional Theory (DFT). DFT is the most commonly used method for finding the ground state energy of interacting electrons, and the most successful one. The central idea is to reduce the complicated many-particle wavefunction to the single-electron density, and encode all the interaction terms in a so-called "universal functional". This is motivated by the observation that the only variable term in the problem is the external potential which only depends on the single-electron density, while both the kinetic (i.e. phase-sensitive) and the interaction term are fixed, making the functional independent of the instance of the problem and thus "universal". The central task in DFT is now to find the universal functional, or a good approximation thereof. However, as we will show, the existence of an efficiently computable universal functional – or even of polynomial accuracy approximations – allows to place INTERACTING ELECTRONS in NP, and would thus imply NP = QMA which is both considered unlikely and would resolve long-standing open problems, explaining the difficulty – or impossibility – of finding the universal functional.

Finally, we turn towards Hartree-Fock, a precursor of DFT and still widely used in numerical calculations. Different from DFT, it is based on an ansatz for the wave function and is therefore not guaranteed to give the optimal ground state, while on the other hand it can be applied without having to guess some functional. We show that these differences also affect the computational complexity of the problem and prove that finding the minimum in Hartree-Fock is NP-complete.

Our proofs will make excessive use of perturbation theory constructions. Rigorous derivations of perturbation expansions suitable for our purposes can be found in [9, 13], and we tacitly use these techniques in our perturbation expansions.

2 Reduction to a spin model

The first part of proving the QMA-hardness of INTERACTING ELECTRONS consists of reducing the complicated continuous three-dimensional many-electron system to a two-dimensional square lattice of qubits interacting via a fixed exchange interaction and subject to a tunable local field.

We start by going to the two-dimensional Hubbard model [6], which describes a system of Fermions hopping on a lattice. It is a well-known special case of the Hamiltonian (1), obtained by putting a strong trapping field in the z direction and a periodic lattice potential in the x-y direction. This confines the electrons to localized fermionic modes on a two-dimensional lattice, described by fermionic creation operators $a^\dagger_{i,s}$. Here, i denotes the site on the 2D lattice, and $s \in \{0,1\}$ the spin state.

The low-energy sector of the resulting system is described by the effective Hamiltonian

$$H_{\mathrm{Hubbard}} = -t \sum_{<i,j>,s} a^\dagger_{i,s} a_{j,s} + U \sum_i n_{i,0} n_{i,1} - \sum_i \vec{B}_i \cdot \vec{\sigma}_i \ , \tag{2}$$

where $<i,j>$ denotes nearest neighbors on the 2D lattice, $n = a^\dagger a$ is the particle number operator, and $\vec{\sigma}_i$ is the vector of Pauli matrices acting on the spin degree of freedom at site i. The first term in the Hamiltonian describes an electron tunneling from one site to the adjacent one without changing its spin; its strength is governed by the height and the width of the potential barrier. The second term gives the on-site repulsion: although each site only allows for one spatial mode, this mode can be occupied by one spin up and one spin down electron in a singlet state. However, in this case the energy is increased due to the Coulomb repulsion between the electrons, giving rise to the on-site repulsion U; its strength depends on the width of the potential well. Finally, the rightmost term contains the contribution from the magnetic field which imposes a local field on the spin at each site i. Note that this is the only term which we need to tune locally.

The Hubbard Hamiltonian is thus obtained from the perturbation expansion of the low-energy spectrum of interacting electrons in a particular external potential; a detailed derivation, including bounds on the higher-order terms, can be found e.g. in [2, 8]. Except for being a special case of INTERACTING ELECTRONS, the Hubbard model is of high interest on its own, as it is believed to be the appropriate model for understanding effects like high-temperature superconductivity, quantum magnetism, and heavy fermions. Indeed, it is one of the most intensively studied models in solid state physics, making the investigation of its computational complexity interesting on its own. A particular nice thing about the Hubbard model is that deciding its ground state energy is clearly inside QMA, since it can be mapped to a spin system via the Jordan-Wigner transform. Although this makes the Hamiltonian non-local, its ground state energy can still be measured efficiently [12], and with what follows, we have altogether that the Hubbard model is QMA-complete.

To show the QMA-hardness of the problem, we transform the Hubbard Hamiltonian Eq. (2) to a Hamiltonian on a 2D square lattice of spins. Therefore, we choose the on-site repulsion energy U very large as compared to t (explicit choices for the scalings are given later), and operate the system in the so-called half-occupancy regime, i.e. where there are as many electrons as sites. Due to the large on-site repulsion, the hopping is suppressed to order t/U, and in the ground state each mode will be occupied by exactly one electron, providing a spin degree of freedom which we can use as a qubit. The coupling between the spins is obtained in second-order perturbation theory and is due to a virtual process where one spin tunnels to the adjacent site, interacts with the other spin, and tunnels back. Since two spins can only be at the same site if they form a singlet, this process decreases the energy of a singlet state by $4t^2/U$, and since $(\mathbb{1} - \vec{\sigma}_i\vec{\sigma}_j)/4$ is the projector onto the singlet, the effective interaction Hamiltonian is

$$H_{\mathrm{exc}} = \frac{t^2}{U}\vec{\sigma}_i \cdot \vec{\sigma}_j - \frac{t^2}{U} + O(t^3/U^2) \ .$$

Omitting the energy offset, we are left with an effective lattice Hamiltonian

$$H_{\mathrm{Heis}} = \frac{t^2}{U}\sum_{<i,j>} \vec{\sigma}_i \cdot \vec{\sigma}_j - \sum_i \vec{B}_i \cdot \vec{\sigma}_i + O(Nt^3/U^2) \ . \tag{3}$$

3 QMA-completeness of H_{heis}

In the following, we show that deciding the ground state energy of the Hamiltonian H_{Heis} with arbitrary local field is a QMA-complete problem. Note that it is clearly in QMA, as it is a special instance of the LOCAL HAMILTONIAN problem [10]. To prove QMA-hardness, we show that the problem can be reduced to the ground state problem of [13], where QMA-completeness is shown for a Hamiltonian on a 2D square lattice with interactions of the form

$$H_{2D} = \sum_{<i,j>} \lambda_{ij} A_{(ij)} \otimes B_{(ij)} \ ,$$

where $A_{(ij)}$ and $B_{(ij)}$ are Pauli matrices. Following [13], we call interactions of the form $\lambda_{ij} A_{(ij)} \otimes B_{(ij)}$ Pauli interactions.

We will now show how a 2D lattice with arbitrary Pauli interactions can be embedded into a 2D lattice with fixed Heisenberg interactions and local fields, Eq. (3). Therefore, we will use a number of so-called perturbation gadgets. These basic building blocks consist of a few qubits where a strong local field is applied to one of them, so that the interactions of the Hamiltonian induce an effective coupling on the remaining qubits in the low-energy subsector. An important observation made in [13] is that these gadgets can be combined: firstly, one can use several gadgets in parallel, as long as they do not share an interaction term, and secondly, one can stack gadgets atop of each other, as long as all new terms together are still in the perturbation regime, i.e. much smaller than the strong fields of the previous gadgets. In particular, this can be accomplished with polynomial-strength fields if the number of "gadget layers" is kept constant. For

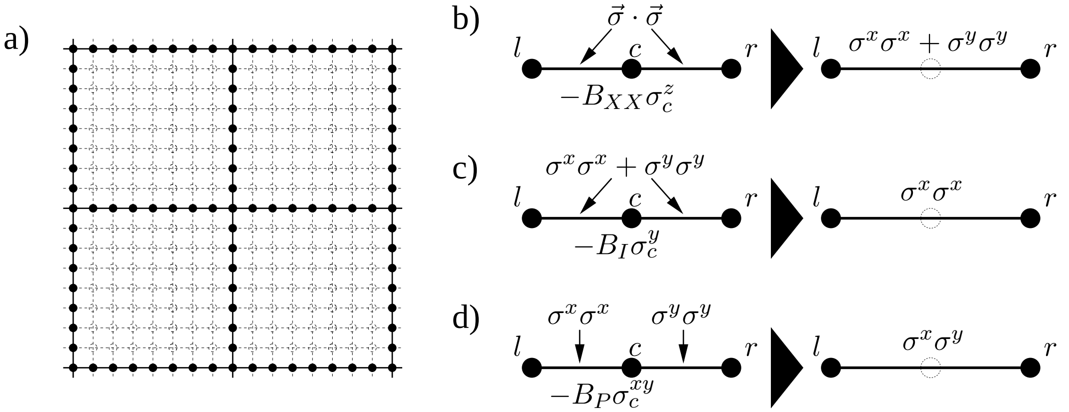

Figure 1: Gadgets used to go from the 2D lattice with fixed Heisenberg interactions and local field to the 2D lattice with tunable Pauli interactions.

the moment, we will discuss our gadgets one after another and postpone the discussion on how to appropriately connect them.

We will use a total of 4 different gadgets which we apply atop of each other. The first gadget erases vertices from the 2D lattice, and we use it to create a sparse lattice where each effective edge consists of 8 edges with a Heisenberg coupling, see Fig. 1a. The second gadget reduces the full Heisenberg coupling to an XX coupling, thus mapping two edges to one (Fig. 1b). The third gadget reduces two XX edges to one with an Ising-type coupling (Fig. 1c), and the last gadget combines two Ising edges to obtain an arbitrary Pauli coupling of variable strength (Fig. 1d). By properly concatenating the latter three, one can implement any Pauli coupling on each effective edge of the sparse lattice. The following derivations will make use of the techniques for perturbation gadgets as found in [9, 13].

The first gadget erases a node from the graph. Therefore, we apply a strong field B_e to this qubit in, say, the z direction. The Hamiltonian is

$$H_{\text{erase}} = -B_e\sigma_c^z + \lambda \sum_{q=l,r,t,b} \vec{\sigma}_c \cdot \vec{\sigma}_q \ ,$$

where in the following $\lambda = t^2/U$ [cf. Eq. (3)], and where c, l, r, t, and b denote the center, left, right, top, and bottom qubit, respectively. Doing a perturbation expansion in λ/B_e, we find

$$H_{\text{erase}} = \left[-B_e + \lambda \sum_{q=l,r,t,b} \sigma_q^z + O\left(\frac{\lambda^2}{B_e}\right)\right] \otimes |0\rangle\langle 0|_c \ .$$

Thus, the gadget induces an effective local Hamiltonian on the adjacent sites while decoupling them up to a perturbation of order λ^2/B_e. Since the induced local terms can be easily compensated by changing the corresponding local fields, the effect of this gadget is to erase the central qubit for sufficiently large B_e.

We now use the erasure gadget to cut a sublattice into the square lattice. Therefore, we apply the field B_e to all sites we want to remove, so that we are left with a sparse lattice as in Fig. 1a, where each edge contains 8 edges of the original lattice. These edges will be used to build a single edge with an arbitrary Pauli coupling, using a sequence of gadgets each of which maps two edges to one while creating a new interaction. The first of these gadgets reduces the Heisenberg interaction to an XX interaction, i.e. an interaction of the form $\sigma_i^x\sigma_j^x + \sigma_i^y\sigma_j^y$. Therefore, consider three qubits on a line and apply a field B_{XX} to the central qubit,

$$H_{XX} = -B_{XX}\sigma_c^z + \lambda \sum_{q=l,r} \vec{\sigma}_c \cdot \vec{\sigma}_q \ .$$

Doing the perturbation expansion to second order, we obtain

$$H_{XX} = \left[H_{\mathrm{loc}} - \frac{\lambda^2}{4B_{XX}} \left(\sigma_l^x\sigma_r^x + \sigma_l^y\sigma_r^y\right) + O\left(\frac{\lambda^3}{B_{XX}^2}\right)\right] \otimes |0\rangle\langle 0|_c \ ,$$

where $H_{\mathrm{loc}} = -B_{XX} + \lambda \sum_{q=l,r} \left(Z_q + (\lambda/4B_{XX})(Z_q - \mathbb{1})\right)$ collects all local terms, which can be compensated easily. Clearly, by choosing a different direction for B_{XX}, we can equally get couplings in x-z or in y-z.

The next gadget reduces the XX coupling to an Ising type coupling $\sigma_l^x\sigma_r^x$. Therefore, we take a chain of three qubits coupled by the XX interaction (w.l.o.g. in x and z direction), and apply a field in Z direction to the second qubit,

$$H_{\mathrm{Ising}} = -B_I\sigma_c^z - \mu \sum_{q=l,r} \left(\sigma_l^x\sigma_r^x + \sigma_l^z\sigma_r^z\right)$$

with $\mu = \lambda^2/4B_{XX}$, which perturbatively gives

$$H_{\mathrm{Ising}} = \left[-B_I - \mu \sum_{q=l,r} Z_q - \frac{\mu^2}{B_I}(\sigma_l^x\sigma_r^x + \mathbb{1}) + O\left(\frac{\mu^3}{B_I^2}\right)\right] \otimes |0\rangle\langle 0|_c$$

and is thus locally equivalent to an Ising coupling.

The last gadget is being used to get an arbitrary Pauli interaction with tunable strength. In order to get e.g. a coupling of the form $\sigma_l^x\sigma_r^y$, we take a three-qubit chain with two different Ising couplings and apply a strong diagonal field on the central qubit,

$$H_{\mathrm{Pauli}} = -B_P\sigma_c^{xy} - \vartheta(\sigma_l^x\sigma_c^x + \sigma_c^y\sigma_r^y) \ ,$$

with $\vartheta = \mu^2/B_I$, and where $\sigma^{xy} = |+\rangle_{xy}\langle +|_{xy} - |-\rangle_{xy}\langle -|_{xy}$ with $|\pm\rangle_{xy} = (|0\rangle + \sqrt{i}|1\rangle)/\sqrt{2}$, $\sqrt{i} = (1+i)/\sqrt{2}$ vectors along the $x+y$ direction. The resulting interaction is is obtained from the perturbation series as

$$H_{\mathrm{Pauli}} = \left[-B_P - \frac{\vartheta}{\sqrt{2}}(\sigma_l^x + \sigma_r^y) - \frac{\vartheta^2}{2B_P}(\sigma_l^x\sigma_r^y + \mathbb{1}) + O\left(\frac{\vartheta^3}{B_P^2}\right)\right] \ , \tag{4}$$

thus enabling us to construct any desired Pauli coupling. Note that this construction also allows to keep an Ising interaction $\sigma_l^x\sigma_r^x$ by applying B_P in x direction.

The strength of the Pauli couplings can be tuned by appropriately changing the corresponding field B_P. However, there are a two caveats: On the one side, we want the $O(\vartheta^3/B_P^2)$ term above to vanish polynomially (in particular faster than the desired accuracy in the ground state energy) and are thus restricted to interaction strengths $\vartheta/B_P \lesssim 1/\mathrm{poly}(N)$. However, since the interaction strength in the QMA-complete construction for 2D lattices with Pauli interactions in [13] is polynomially bounded, we can rescale all interactions as well as the desired energy precision by a polynomial.

On the other side, in order to obtain very weak couplings we would need to choose B_P arbitrarily large. Since, however, we want to keep all our Hamiltonian terms bounded by some polynomial in the systems size N, we cannot obtain couplings smaller than $1/\mathrm{poly}(N)$. Again, this can be resolved by setting these coupling terms to zero, which results in an overall error of at most $N/\mathrm{poly}(N)$ in energy. By choosing this sufficiently smaller than the desired precision, we can also avoid problems in this case. Clearly, setting the interaction terms to zero can be done by erasing those edges right in the first step when the erasure gadget is applied.

Let us now go through the whole procedure of rescaling the interactions, choosing the local field strength, and assembling the gadgets, as to see that we stay in a polynomial regime. Recall that that when assembling the gadgets, the energy scale of all newly added gadgets should still be small compared to the large field underlying the perturbation expansion in the previous stage. We start from a 2D lattice with arbitrary Pauli interactions, where the interaction strength is bounded by a polynomial $a(N)$, and with a desired energy accuracy $1/b(N)$ with b a polynomial. Now choose a polynomial $c \gg a^2 b$ and rescale everything by $1/c$, thus obtaining a maximum interaction strength a/c and a precision $1/bc$.[1] By omitting all interactions weaker than $O(1/N^2 b(N)c(N))$, we make an error of $O(1/Nb(N)c(N))$ in energy which is vanishing relative to the precision.

We now let $U = 1$ and $t = O(1/\sqrt{bc})$ in the Hubbard model (2), leading to $\lambda = O(1/bc)$. Then, we choose the following magnetic fields:

$$\begin{aligned} B_e &= O\left(N^6\lambda(bc)^{2/7}\right) \ , \\ B_{XX} &= O\left(N^4\lambda(bc)^{2/7}\right) \ , \\ B_I &= O\left(N^2\lambda(bc)^{2/7}\right) \ , \\ B_P &= O\left(\lambda(bc)^{2/7}\right) \ , \end{aligned}$$

which ensures four things: firstly, each field is much stronger than all following contributions, second, each field is much stronger than the corresponding perturbation λ, μ, or ϑ, thirdly, the permissible interaction strength $\vartheta^2/2B_P$ is large enough to allow for $O(1/N^2 b(N)c(N))$ couplings, and finally, all the error terms from the perturbation expansions are $O(N/(bc)^{9/7})$ and thus for sufficiently large c much smaller than the desired precision which is $O(1/bc)$.

---

[1] For convenience, we omit the parameter N when appropriate.

Thus, by assembling all the gadgets together we finally find that INTERACTING ELECTRONS is QMA-hard.

4 Density Functional Theory

As we have seen, finding the ground state energy of systems of interacting electrons is both central in physics and a computationally hard problem. Therefore, approximation methods have to be used in order to obtain results for realistic scenarios. One of the most popular methods for this problem is Density Functional Theory (DFT) [5, 14]. It is based on the observation that the only varying part of the Hamiltonian (1) is the potential term V which only depends on the single electron density, while all the other terms, and in particular the interactions, are fixed. Therefore, one can split the minimization over the set of antisymmetric N-particle wavefunctions Ψ in two parts: One considers the one-particle reduced density matrix given by

$$\rho(r) = \sum_{s_1, s_1', s_2, \dots, s_N} \int_{\mathbb{R}^{3N-3}} \mathrm{d}r_2 \cdots \mathrm{d}r_N \Psi(r_1, s_1, \dots, r_N, s_N) \times \\ \Psi^*(r_1, s_1', r_2, s_2, \dots, r_N, s_N) |s_1\rangle\langle s_1'|$$

which already fully determines the energy contribution coming from the the potential V. (Note that we can replace the spin density matrix by a vector in the Bloch sphere, making the symmetry between spatial and spin density more clear.) On the other hand, for a given $\rho(r)$ one can minimize the energy of the fixed terms $T + I$ (the kinetic and the interaction term) over all N-electron states compatible with the marginal $\rho(r)$ without having knowledge about V. Thus, the ground state energy of the problem is given by

$$E_0 = \min_{\rho} \{ \mathrm{tr}(V\rho) + F_N[\rho] \} \ , \tag{5}$$

where the "universal functional"

$$F_N[\rho] = \min_{\Psi \to \rho} \left(\int \Psi^* (T + I) \Psi \right)$$

minimizes the fixed kinetic and interaction term over all N-electron states compatible with ρ. Therefore, the central problem in DFT is to obtain the functional F_N, or at least a good enough approximation for the setup under investigation.

In the following, we show that one cannot hope to obtain an efficiently computable universal functional F_N. Therefore, consider again the 2D Hubbard model which we know to be QMA-complete, and let us show that we can specify $\rho(r)$ for the ground state classically.[2] Therefore, note that we can specify the probability distribution over the discrete set of N lattice sites efficiently, and since the wavefunctions corresponding to the localized modes $a_{i,s}^\dagger$ are known, this suffices to reconstruct the single-electron density matrix $\rho(r)$. Thus, if F_N could be computed efficiently, one could provide an efficiently checkable witness for the ground state, and thus the ground state energy of the Hubbard model could be decided in NP, implying NP=QMA. Therefore, there cannot be an efficiently computable universal functional unless NP=QMA, which is considered unlikely. However, even if one is willing to accept NP=QMA, our result still emphasizes the hardness of finding the functional, since $\mathsf{NP} = \mathsf{QMA} \Rightarrow \mathsf{NP} = \mathsf{MA} \Rightarrow \mathsf{NEXP} \neq \mathsf{MA} \Leftrightarrow \mathsf{NEXP} \not\subset \mathsf{P/poly}$ [7], and whether NEXP contains hard Boolean functions is a long-standing open problem in computational complexity.

Let us note that the argument for the impossibility of finding a universal functional also holds for approximations up to polynomial precision: for everything we have to know in the definition of QMA is the ground state energy up to a polynomial accuracy.

One of the precursors of DFT which is still widely used is the Hartree-Fock method [11]. It is similar to DFT in that it reduces the N-electron equation to a problem of individual electrons moving in an external field which depends only on the average electron distribution. Different from DFT, Hartree-Fock is based on a particular ansatz for the wave function and is therefore not guaranteed to give the true ground state energy; on the other hand, it is an iterative method which can be applied without prior assumptions, whereas in DFT some *a priori* guess on the form of the universal functional has to be made.

The starting point is a Hamiltonian

$$\mathcal{H} = \sum H^{(1)}_{i,j} a_i^\dagger a_j + \sum H^{(2)}_{ij,kl} a_i^\dagger a_j^\dagger a_k a_l + \dots \ , \qquad (6)$$

which is local, meaning the series terminates at some fixed $H^{(r)}$, and with a number of fermionic modes $M \geq N$; if it is derived from two-particle interactions as in the Schrödinger equation (1), $r = 2$. The Hartree-Fock method approximates the ground state of this Hamiltonian using the ansatz $b_N^\dagger \cdots b_1^\dagger |\Omega\rangle$ with $b_i = \sum u_{ij} a_j$ (where $|\Omega\rangle$ is the vacuum). Note that this corresponds to an antisymmetrized product of single-particle wave functions, which is how Hartree-Fock is usually presented.

In the following, we show that approximating the ground state energy using the Hartree-Fock method is an NP-complete problem. More precisely, we consider the problem of deciding whether the lowest energy of (6) within the Hartree-Fock ansatz is below some a or above some $b > a$. We show that the problem is inside NP for up to an exponential accuracy $b - a$ and for any r, and that for NP-completeness a polynomial accuracy $b - a < 1/\mathrm{poly}(N)$ and $r = 2$ are sufficient.

---

[2] Note that this does not place the problem in QCMA [10] since we do not know how to obtain the optimal Ψ from ρ even quantumly.

To see that the problem is in NP, note that a Hartree-Fock state is fully characterized by the u_{ij}'s, and that from there its energy can be computed efficiently. Conversely, the problem is shown to be NP-hard by mapping it to the ground state problem for Ising spin glasses which is known to be NP-hard [3]: Given an $L\times L\times 2$ lattice of two-level spins $S_i = \pm 1$ with a nearest neighbor Ising coupling $\mathcal{H} = \sum J_{ij}S_iS_j$, $J_{ij} \in \{0, -1, 1\}$, determine whether the ground state energy is the minimum one allowed by the individual J_{ij}'s or not. Therefore, embed the $N = 2L^2$ classical spins into a fermionic system with $2N$ modes occupied by N fermions. The modes come in pairs (a_{2i}, a_{2i+1}), and a Hamiltonian term $\lambda n_{2i}n_{2i+1}$, $\lambda = O(N^2)$ penalizes double occupancy, so that in the ground state exactly one mode per pair is occupied, giving an effective spin degree of freedom [12]; the coupling $J_{ij}S_iS_j$ of these spins is realized as $J_{ij}\sum_{p,q=0,1}(-1)^{p+q}n_{2i+p}n_{2j+q}$. As the ground state of the system is a classical spin state, it can be expressed as a Hartree-Fock state where $b_i = a_{2i}$ or $b_i = a_{2i+1}$, respectively, and since the classical Hamiltonian has a constant gap while perturbations from the penalized subspace are at most $O(1/\lambda^2)$, a polynomial accuracy is sufficient to make the problem NP-hard.

5 Outlook

We have shown that INTERACTING ELECTRONS, one of the core problems in physics and theoretical chemistry, is QMA-hard. On our way, we have proven QMA-completeness both for the Hubbard model and for qubits in a 2D lattice with a fixed Heisenberg coupling and local fields. We have then applied our findings to Density Functional Theory to show that the existence of an efficiently computable universal functional would place the Hubbard model and thus QMA in NP, making the existence of an explicit functional unlikely. Finally, we have shown that solving a system of interacting electrons using Hartree-Fock is NP-complete.

An open question is whether INTERACTING ELECTRONS can be placed inside QMA. The difficulty comes from the continuous nature of of the wavefunction, as compared to the Hubbard model where one can provide a proof in second quantization. Therefore, one would have to discretize the wavefunction in order to be able to encode the proof in a finite-size quantum state. In principle this should be feasible, since the kinetic term penalizes fast fluctuations in the wave function, thus justifying a coarse graining of the ground state witness. However, care has to be taken concerning the magnetic field since fast changes of the spin are not detected by the kinetic term. Therefore, in order to place INTERACTING ELECTRONS in QMA, one will have to put constraints on the magnetic field bounding the speed of possible fluctuations.

Another open question concerns the computational complexity of the problem of V-representability [14]: given a local density $\rho(r)$, does there exist a potential $V(r)$ such that the Hamiltonian $T + I + V$ [Eq. (1)] has a unique ground state with local density $\rho(r)$? The lattice version of the problem is the following: Given a local density ρ on the lattice, decide whether there exists a potential V such that $H = H_0 + V$ (for a fixed H_0) has a unique ground state with density ρ. On the one hand, it can be proven in QMA whether a given density is the ground state to some local (spin or fermionic) Hamiltonian H – Merlin provides H, the ground state energy E_0, and an associated (possibly mixed) ground state γ, and Arthur either checks that γ is the ground state of the local Hamiltonian $H' = (H - E_0)^2$, or that it indeed gives rise to the proper reduced density. However, it is unclear how to simultaneously prove *uniqueness* of the ground state – it rather seems that *dis*proving the uniqueness of the ground state for a given Hamiltonian is in QMA (Arthur provides two non-identical ground states). Note that without the requirement of uniqueness of the ground state – i.e. asking whether there exists a potential V such that some (mixed) ground state γ gives rise to ρ – the lattice version of V-representability can probably be decided efficiently [4].

Acknowlegdements

We thank H. Buhrman, I. Cirac, G. Giedke, J. Kempe, and B. Toner for helpful discussions and comments. N.S. thanks the Erwin Schrödinger Institute in Vienna, where a major part of this work was carried out, for their hospitality.

References

[1] D. Aharonov, D. Gottesman, and J. Kempe, arXiv:0705.4077 (2007); S. Irani, arXiv:0705.4067 (2007).

[2] A. Auerbach, *Interacting Electrons and Quantum Magnetism*, (Springer, New York, 1994).

[3] F. Barahona, J. Phys. A **15**, 3241 (1982).

[4] J.T. Chayes, L. Chayes, and M.B. Ruskai, J. Stat. Phys. **38**, 497 (1985).

[5] P. Hohenberg and W. Kohn, Phys. Rev. **136**, B864 (1964); W. Kohn and L.J. Sham, Phys. Rev. **140**, A1133 (1965).

[6] J. Hubbard, Proc. R. Soc. Lond. A **276**, 238 (1963); J. Hubbard, Proc. R. Soc. Lond. A **277**, 237 (1964).

[7] R. Impagliazzo, V. Kabanets, and A. Wigderson, Proceedings of the Sixteenth IEEE Conference on Computational Complexity, 1 (2001).

[8] D. Jaksch, *Bose-Einstein condensation and Applications*, PhD thesis (1999).

[9] J. Kempe, A. Kitaev, and O. Regev, SIAM J. Comp., **35**, 1070 (2006); quant-ph/0406180 (2004).

[10] A.Y. Kitaev, A.H. Shen, and M.N. Vyalyi, *Classical and quantum computation* (American Mathematical Society, Providence, Rhode Island, 2002); D. Aharonov and T. Naveh, quant-ph/0210077 (2002).

[11] I.N. Levine, *Quantum chemistry* (Prentice-Hall, Upper Saddle River, New Jersey, 2000).

[12] Y.-K. Liu, M. Christandl, and F. Verstraete, Phys. Rev. Lett. **98**, 110503 (2007); quant-ph/0609125.

[13] R. Oliveira and B.M. Terhal, quant-ph/0504050 (2005).

[14] R.G. Parr and W. Yang, *Density-Functional Theory of Atoms and Molecules*, (Oxford University Press, Oxford, 1989); R.M. Dreizler and E.K.U. Gross, *Density Functional Theory* (Springer, New York, 1990).